AD-A103 130 THE THEORY AND APPLICATION OF DECISION ANALYSIS(U) 1/1
STANFORD UNIV CA R A HOWARD 31 JAN 80 N00014-79-C-0036

UNCLASSIFIED F/G 12/4 NL

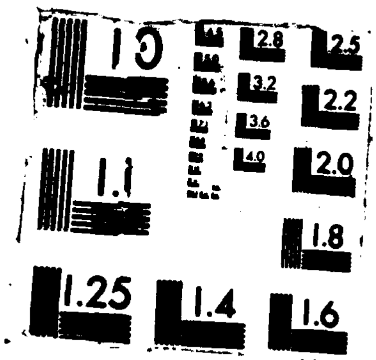
1.0
2.8
2.5
3.2
2.2
3.6
4.0
2.0
1.1
1.8
1.25
1.4
1.6

FINAL TECHNICAL REPORT

to

ADVANCED RESEARCH PROJECTS AGENCY

Washington, D.C.

on 197-054

A Program of Research on

THE THEORY AND APPLICATION OF DECISION ANALYSIS

Ronald A. Howard, Principal Investigator

(415) 497-4176

Contract Number: N00014-79-C-0036

Amount of Contract: $49,990

Effective Date of Contract: October 1, 1978

Contract Expiration Date: September 30, 1979 (with extension granted of three months' time, thereby expiring on December 31, 1979

Date of Report: January 31, 1980

Contract Period Covered: 10/1/78 - 12/31/79

Work Unit: NR 197-054

Scientific Officer:

Director, Engineering Psychology Programs
Office of Naval Research
Department of the Navy
800 North Quincy Street
Arlington, Virginia 22217

Name of Contractor:

Board of Trustees of the
Leland Stanford Junior University
c/o Office of Research Administrator
Encina Hall
Stanford, California 94305

RESEARCH ACCOMPLISHMENTS

Research was conducted in five main areas that were at least partially supported by this contract.

Information Perishing

The principal topic investigated was information perishing. This research was reported in Research Report No. EES-DA-79-1, "Optimal Information Acquisition for Randomly Occurring Decisions", by Ali M. Sharifnia, dated September 1979, described under the contract. The work is well described by the abstract of this report:

> "This research investigates the information acquisition policies for randomly occurring decisions. A randomly occurring decision is a decision that must be made upon the occurrence of a precipitating event that occurs randomly with time. Due to the urgency often associated with this type of decision, it is difficult to obtain fresh information at the time of the decision. Therefore, the available information will be limited to the unfresh (old) information from the past. Since the information becomes outdated and obsolete in time, regular updating of the information is often desirable in order to be prepared for the decision.
>
> The process of information outdating (perishing) and its relationship with the characteristics of the dynamic environment, the decision for which the information is used, and the type of information (perfect or imperfect) are investigated. Based on this process, as well as the cost of information recovery (updating), and the likelihood of the occurrence of the decision, the optimal policies for the recovery of information are studied. Policies which use only the prior knowledge about the environment as well as those which utilize the information in each observation (in addition to the prior knowledge) are analyzed, and optimality conditions are found for each case. The case of a one-time decision, namely when the decision happens only once, is studied initially. The results are then extended to the case where the decision may be repeated in time."

A-1

Hazard Decision Analysis

A multi-year research effort on making decisions involving risks of life and death was advanced considerably by this contract and resulted in the publication of the report, "Life and Death Decision Analysis", Research Report No. EES DA-79-2, by Ronald A. Howard, dated December 1979, which was published under under this contract. The abstract describes this work:

> "No assertion can command attention in time of emergency like, 'It's a matter of life and death'. The problem of making decisions that can affect the likelihood of death is one of the most perlexing facing the analyst. As individuals, we are often called upon to make decisions that affect our safety, and others are increasingly making those decisions on our behalf. Yet most present approaches to life and death-decision making concentrate on the value of an individual's life to others rather than to himself. These approaches are both technically and ethically questionable.
>
> In this report, we develop a model for an individual who wishes to make life and death decisions on his own behalf or who wishes to delegate them to his agents. We show that an individual can use this model if he is willing to trade between the quality and the quantity of his life. A simplified version requires him to establish preference between the resources he disposes during his lifetime and the length of it, to establish probability assessments on these quantities, to characterize his ability to turn present cash into future income, and to specify his risk attitude. We can use this model to determine both what an individual would have to be paid to assume a given risk and what he would pay to avoid a given risk. The risks may range from those that are virtually infinitesimal to those that are imminently life threatening. We show that this model resolves a paradox posed by previously proposed models. In this model there is no inconsistency between an individual's refusing any amount of money, however large, to incur a large enough risk, and yet being willing to pay only a finite amount, his current wealth, to avoid certain death.
>
> We find that in the normal range of safety decisions, say 10^{-3} or less probability of death, the individual has a small-risk value of life that he may use in the expected value sense for making safety decisions. This small-risk life value applies both to risk increasing and risk decreasing decisions, and is of the order of a few million dollars in the cases we have measured. This small-risk value of life is typically many times the economic value of life that has been computed by

other methods. To the extent such economic values are used in decisions affecting the individual, they result in life risks that are in excess of what he would willingly accept. Using the small-risk life value as a basis for compensation should allow most risk-imposing projects to proceed without violating anyone's right to be free from significant involuntarily imposed hazards.

The report demonstrates the use of the model to treat hazards that continue over many years, to determine the size of contributions to saving the lives of others, and to incorporate more precise specifications of consumption-lifetime preferences."

Other Topics

This contract partially supported larger research into the areas of crisis decision analysis, decisions that affect the distant future, and influence diagrams. Abstracts of theses completed in these areas are attached.

PERSONNEL SUPPORTED

Ronald A. Howard, Principal Investigator

Patricia A. Owen, Research Assistant
Daniel L. Owen, Research Assistant
Burke E. Robinson, Research Assistant

PUBLICATIONS

"Life and Death Decision Analysis" by Ronald A. Howard, Research Report EES DA-79-2, December 1979.

"An Assessment of Decision Analysis", Ronald A. Howard, Operations Research, Vol. 28, No. 1, January-February 1980.

"On Making Life and Death Decisions", by Ronald A. Howard, Societal Risk Assessment: How Much Is Enough?, edited by R. C. Schwing and W. A. Albers, Jr., General Motors Research Laboratories, Plenum Press (in press).

"Optimal Information Acquisition for Randomly Occurring Decisions", by Ali M. Sharifnia, Research Report EES DA-79-1, September 1979.

"Crisis Decision Analysis" by Burke E. Robinson, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, June 1979.

"Decisions that Affect Outcomes in the Distant Future", by Patricia A. Owen, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, May 1979.

"The Concept of Influence and Its Use in Structuring Complex Decision Problems", by Daniel Owen, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, June 1979.

Burke E. Robinson
June 1979

CRISIS DECISION ANALYSIS

ABSTRACT

The tenor of our times is crisis. As well as danger, a crisis presents an opportunity for decision. But making decisions in a crisis is difficult because surprises often occur, important values are threatened, and time is limited. To help individuals manage crisis problems more effectively, we provide a modeling procedure, adapted from standard decision analysis methods, that guides the structuring, assessment, and analysis of crisis decisions so as to:

- identify and structure all those, and only those, variables that may change the decision,
- assess probable outcomes quickly and efficiently, and
- indicate, at any moment, the best decision and the value of further modeling.

Descriptive studies, explaining and predicting crisis decision-making behavior, are abundant. Normative studies, prescribing how crisis decisions should be made, are scarce. We expand upon results of previous normative research in decision structuring by developing a series of modules for modeling crisis decisions, as shown in Figure 1.1.

A crisis precludes structuring an entire deterministic model; instead, only the most important alternatives, events, and outcomes are identified and structured. A skeleton of the model is constructed from the structuring modules: template, sequence, influence, tree, preference,

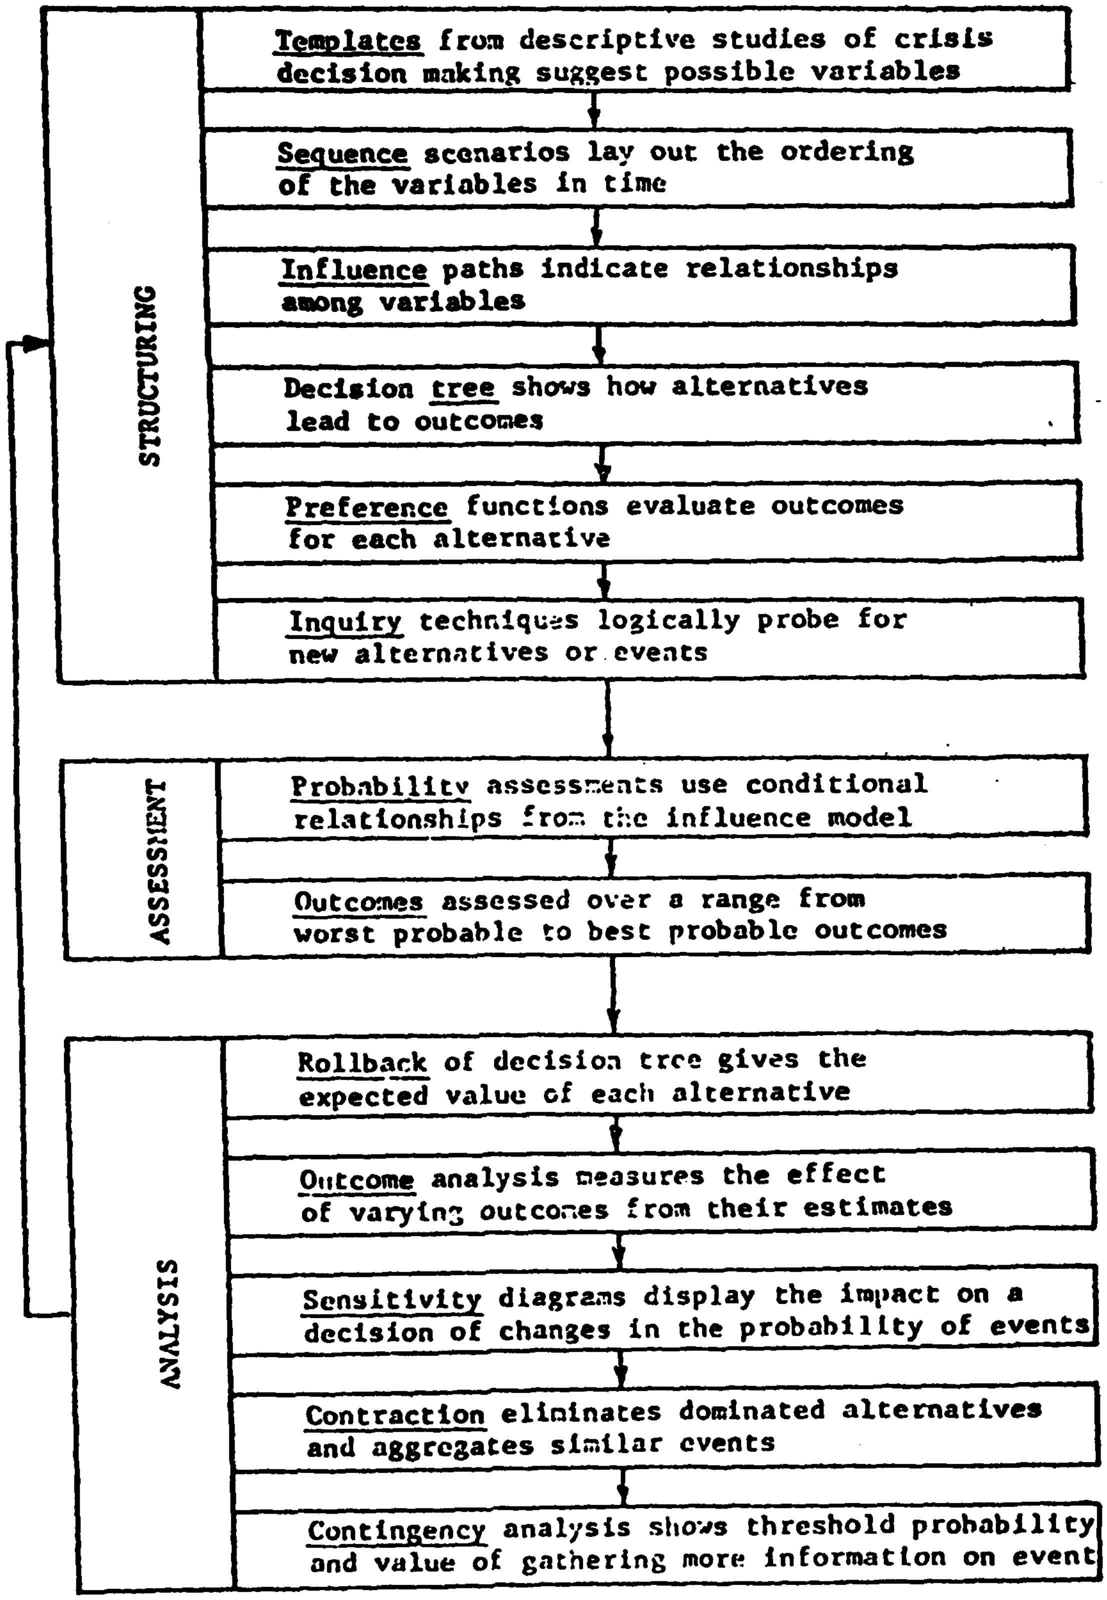

Figure 1.1. Crisis decision analysis procedure

and inquiry. Aids in the template module suggest variables the decision makers may wish to include in their model. The sequence module lays out the order of occurrence of these variables. Paths are then added to indicate the influences of the variables on each other. With minor manipulations, a decision tree can be drawn directly from an influence model. Preference functions evaluate the outcomes for each alternative. Techniques of inquiry are logical questions and probes to identify new events or alternatives.

A crisis also precludes detailed assessments of event probabilities and outcomes; instead, estimates are elicited quickly and efficiently in the modules for probability assessment and outcome assessment. These assessments are made only after other modeling steps identify an unmodeled event or alternative that may change the decision. Probabilities are assessed using the conditional relationships in the influence model. Outcomes are assessed over a range from the worst probable to the best probable outcomes.

Finally, a crisis precludes the complete analysis of a decision problem; instead, at any moment, the best decision and the value of further modeling are indicated from the analysis modules: rollback, outcome, sensitivity, contraction, and contingency. Decision trees are rolled back to find the expected value (or utility) of each alternative. Outcome analysis then measures the effect of varying outcomes from their expected estimates. Sensitivity analysis determines the impact on a decision of changes in the probability of an unmodeled event. Contraction of alternatives occurs as dominated ones are eliminated; contraction of events occurs as they are aggregated by similar effect on a decision. Contingency analysis, after no further modeling is desired,

shows the threshold probabilities and the value of gathering more information for important events.

The facts of the Mayaguez crisis of 1975 provide a realistic setting for a hypothetical example of how the crisis decision analysis procedure would be applied. In an illustrative dialogue between the decision maker, his experts, and a decision analyst, the crisis decision model emerges as the modules are successively invoked.

Another example, a corporate crisis concerning the purchase of ore, shows the application of the modeling procedure to a problem with more specific data. The decision model is developed in phases until all of the important variables are included.

In summary, the crisis decision analysis procedure clearly and concisely guides the modeling of these crisis decision examples. Limitations of the methodology point out areas where further research is needed before the procedure can be used successfully in actual crises.

Patricia A. Owen
May 1979

DECISIONS THAT AFFECT OUTCOMES IN THE DISTANT FUTURE

ABSTRACT

The question that this research addresses is how decisions involving many citizens should be made when those decisions affect outcomes in the distant future. "Distant" means beyond the lifetimes of individuals alive now. The decision maker might be either a private company or a public agency. The contribution of this research is a comprehensive methodology for decision making in this situation, including a mathematical theory and techniques for assessing the required information.

The approach that is taken is a synthesis of concepts from economics with techniques for handling time preferences and uncertain outcomes from the theory of decision analysis. The result is a methodology for deciding whether to accept or reject individual projects with uncertain outcomes on future generations. The fundamental basis for decision making is the amount that current citizens are willing to pay for outcomes accruing to other individuals as well as outcomes affecting their own consumption. One important product of this research is a set of equations for approximating the amounts individuals would pay for a project. The expressions include a wide range of realistic cases, such as non-expected-value preferences, uncertainty in individuals' lifetimes, and outcomes accruing to "others" in the same generation as well as others in the future.

It is shown that, under certain circumstances, an individual's willingness to pay is equal to his consumer surplus. Thus, cost-benefit analysis is a special case of the results derived in this research. However, there is an important philosophical difference between our approach and traditional analysis. The methodology proposed in this research assumes that only the preferences of <u>current</u> citizens enter the decision-making

process. The future counts only to the extent that current individuals decide to value it.

If we accept this assumption, then the "social discounting" technique used in cost-benefit analysis is not an appropriate way to make decisions affecting outcomes on future generations. This is shown using an example of the government decision to store helium underground. In order for cost-benefit analysis to value projects in a manner consistent with current individuals' preferences, the discount rate would have to vary with the distribution of outcomes among people in each generation, how much current individuals value the future, and what it is that is valued about the future. Although we could force the cost-benefit approach to give a consistent answer by using a complicated discount rate, it is more reasonable to base the decision directly on individuals' preferences and the amounts they are willing to pay.

Daniel Owen
June 1979

THE CONCEPT OF INFLUENCE AND ITS USE IN STRUCTURING COMPLEX DECISION PROBLEMS

Abstract

The generality of the decision analysis methodology permits its application to decision problems regardless of the particular discipline or setting in which the problem occurs. Consequently, the decision analyst may be unfamiliar with the relationships of the variables in the problem. One device for communicating those relationships is a diagram identifying the existence of influences between the variables. This research contributes a general mathematical characterization of the influence between random variables. The influence can be characterized by a matrix that is null if and only if no influence exists and otherwise indicates the degree and type of influence by its nonzero elements.

An electrical engineer uses the schematic diagram of a circuit to conceptualize and communicate the relationship between the voltage at different points of an electronic device. The definition of influence can serve the decision analyst in an analogous manner, helping him to conceptualize and communicate the relationship of the probability distributions on different variables in a probabilistic decision model. The definition of influence supports a calculus of influences that allows one to compute the total influence of one variable on another even when there are several intermediate variables. Using this influence calculus, the importance of a particular variable to the decision model can be determined. An immediate consequence is a recommendation for

which variables to include in the model and whether the uncertainty about a variable is important. These recommendations include a new interpretation of deterministic sensitivity.

An important, philosophical result of this research is the demonstration that which variables should be included in the decision model depend on the decision maker's risk attitude. Two decision makers with the same state of information but different risk attitudes should model the same decision problem differently.

Finally, the theoretical basis for the influence definition is different from that of the conventional discretization or decision tree representation for solving decision problems. Since the acceptability of the influence method depends on its accuracy and ease of implementation relative to discretization, the theoretical bases of the influence method and discretization are compared.

END
9-87
DTIC